

A smooth ROC curve estimator based on log-concave density estimates

Kaspar Rufibach
 Division of Biostatistics
 Institute of Social and Preventive Medicine
 University of Zurich
 kaspar.rufibach@ifspm.uzh.ch

January 14, 2013

Abstract

We introduce a new smooth estimator of the ROC curve based on log-concave density estimates of the constituent distributions. We show that our estimate is asymptotically equivalent to the empirical ROC curve if the underlying densities are in fact log-concave. In addition, we empirically show that our proposed estimator exhibits an efficiency gain for finite sample sizes with respect to the standard empirical estimate in various scenarios and that it is only slightly less efficient, if at all, compared to the fully parametric binormal estimate in case the underlying distributions are normal. The estimator is also quite robust against modest deviations from the log-concavity assumption. We show that bootstrap confidence intervals for the value of the ROC curve at a fixed false positive fraction based on the new estimate are on average shorter compared to the approach by Zhou and Qin (2005), while maintaining coverage probability. Computation of our proposed estimate uses the R package **logcondens** that implements univariate log-concave density estimation and can be done very efficiently using only one line of code. These obtained results lead us to advocate our estimate for a wide range of scenarios.

Keywords: Diagnostic test; Log-concave density estimation; Nonparametric estimation; Receiver operating characteristic curve; Sensitivity and specificity

1 Introduction

The receiver operating characteristic (ROC) curve is a common way of assessing the diagnostic accuracy of a diagnostic test with continuous outcome Z that predicts presence or absence of a binary trait, typically a disease. The ROC curve is defined as the plot of the sensitivity (the true positive fraction) against one minus specificity (the false positive fraction) across all possible choices of threshold values θ . Thus, the ROC curve displays the range of possible trade-offs between the true and false positive rates. To

fix ideas, let F and G denote distribution functions of the test result X and Y of a non-diseased and a diseased subject, respectively. If F and G are absolutely continuous, the ROC curve of the test can be expressed as

$$R(t; F, G) = 1 - G(F^{-1}(1 - t)) \quad (1)$$

where F^{-1} is the quantile function of F , and $t \in (0, 1)$ is the false positive fraction corresponding to a cut-off for positivity of the test.

We suppose that a sample $X_1, \dots, X_m \sim F$ from the non-diseased population and a sample $Y_1, \dots, Y_n \sim G$ from the diseased population are available. We further assume that all these observations are mutually independent and denote the empirical distribution functions based on the samples by \mathbb{F}_m and \mathbb{G}_n , respectively. The empirical quantile function for a sample X_1, \dots, X_m is defined as $\mathbb{F}^{-1}(p) = X_i$ if $\mathbb{F}(X_{i-1}) < p \leq \mathbb{F}(X_i)$ for any $p \in [0, 1]$ and by setting $X_0 = -\infty, X_{n+1} = \infty$. In the absence of any further information about F and G , plugging in these empirical distribution functions in (1) yields the nonparametric maximum likelihood estimator $\tilde{R}_{m,n}(t) = R(t; \mathbb{F}_m, \mathbb{G}_n)$ of $R(t; F, G)$, an increasing step function that simply consists of plotting the empirical proportions $\#\{Y_j > t\}/n$ vs. $\#\{X_i > t\}/m$ for varying t . Strong consistency and strong approximation properties of $\tilde{R}_{m,n}$ are provided in Hsieh and Turnbull (1996). Being a fully nonparametric estimator, $\tilde{R}_{m,n}$ displays the data pattern well but may suffer from substantial variability and, as a step function, is not smooth. Moreover, due to the rugged form, an empirical ROC curve can have the same true positive fraction corresponding to a range of false positive fractions. This is inconvenient when one is interested in finding a particular false positive fraction at a specified true positive fraction.

In applications, it seems sensible to assume that the true underlying ROC curve R is in fact smooth. Thus we expect an estimator to perform better if some smoothness is imposed (Lloyd and Yong, 1999, Section 1). The most obvious way of doing so is to assume a parametric model, estimate the distribution functions F and G on the basis of the samples of healthy and diseased subjects, and compute the corresponding ROC curve. Assuming a normal distribution both in the control and in the cases group is the most prominent parametric model (Pepe, 2003, Example 5.3).

Among the frequentist approaches that are between the two “extreme cases”, the entirely nonparametric and the parametric ROC curve estimate, are semiparametric models. Such models do not assume a parametric form for the constituent distributions but rather for the ROC curve directly (Cai and Moskowitz, 2004). The most prominent semiparametric model for the ROC curve stipulates the existence of a function h such that $h(X) \sim N(0, 1)$ and $h(Y) \sim N(\mu, \sigma^2)$. The binormal ROC curve is then given by

$$\bar{R}_{m,n}(t) = \Phi(a + b\Phi^{-1}(t)) \quad (2)$$

where a and b are to be estimated and Φ is the cumulative distribution (CDF) of a standard normal distribution. Hsieh and Turnbull (1996) propose a generalized least squares estimate whereas Cai and Moskowitz (2004) and Zhou and Lin (2008) discuss profile likelihood methods to estimate a and b in the general setup (2).

With the purpose of defining a smooth, though more flexible than parametric, ROC curve estimate Lloyd (1998) proposes to plug in kernel estimates of F and G into (1). As is common for all kernel estimates, a kernel and, more importantly, a bandwidth have to be

chosen. In the ROC curve estimation problem this issue is even more pronounced, since both the ordinate and coordinate of the estimated ROC curve are random and it is thus unclear which variation to minimize in mean square. Lloyd (1998) suggests to circumvent the bandwidth choice problem by a rather involved graphical examination of bias and variation from bootstrap samples. That this kernel ROC estimate is more efficient than the empirical is shown in the subsequent paper Lloyd and Yong (1999). However, as discussed in the bandwidth comparison study performed by Zhou and Harezlak (2002), Lloyd's estimate may not be practical for routine use in medical studies, although it has good theoretical properties. Its main drawback is that bandwidth selection is ad-hoc and difficult. Hall and Hyndman (2003) take a somewhat different approach as they choose the optimal bandwidth not separately for F and G , but they derive optimal bandwidths via minimizing the mean integrated squared error (MISE) for the estimation of R directly. They show via simulations that their estimator improves on existing approaches in terms of MISE.

Another approach to smooth estimation is taken in Qiu and Le (2001) who combine kernel estimation of G with a quantile estimator due to Harrell and Davis (1982). The latter has been shown to perform better than the usual empirical quantile estimator in many scenarios. Du and Tang (2009) propose a monotone spline approach to ensure monotonicity and transformation invariance of the estimated ROC curve. However, here one has to select a smoothing parameter and Du and Tang (2009) propose to use cross validation to guide that choice.

Further papers that deal with semiparametric approaches to estimate a ROC curve are Lloyd (2002), Peng and Zhou (2004), and Wan and Zhang (2007). However, it seems safe to say that diagnostic tests in applications are to a large extent assessed by using the empirical ROC curve estimate $\hat{R}_{m,n}$ (Zhou and Harezlak, 2002, Section 1) and to a lesser extent the binormal estimate, despite the well-known deficiencies of the latter. Presumably, this has to be attributed to lack of easy accessible software that implements alternative methods.

In Section 2 we briefly summarize some facts on log-concave density estimation. Our new estimator is introduced in Section 3 and some theoretical results are provided in Section 4. An illustration of the proposed ROC curve estimate using a real-world dataset is discussed in Section 5 and its performance in simulations is assessed in Section 6. Bootstrap confidence intervals for the value of the ROC curve at a fixed false positive fraction t are introduced in Section 7 and Section 8 discusses the features of the new estimator under misspecification. Finally, in Section 9 we draw some conclusions. Additional results on log-concave density estimation, on estimation of AUC, on simulations and under misspecification as well as proofs are postponed to the appendix.

2 Log-concave density estimation

In Section 3 we propose an alternative ROC curve estimator in the spirit of Lloyd (1998). Namely, we also model the constituent distributions F and G nonparametrically. However, we do not suggest to use kernels, but the log-concave density estimate initially introduced in Walther (2002), Rufibach (2006), Pal et al. (2007), and Rufibach (2007). The features of the new ROC curve estimator are discussed in Section 3. More details

on recent research on log-concave density estimation is provided in the appendix.

Rate of uniform convergence for the univariate log-concave maximum likelihood estimate was derived in Dümbgen and Rufibach (2009) and its computation is described in Rufibach (2007), Dümbgen et al. (2010), and Dümbgen and Rufibach (2011). The corresponding software is available from CRAN as the R package **logcondens** (Rufibach and Dümbgen, 2011). The multivariate case has been treated by Cule et al. (2010) and the extension to discrete observations in Balabdaoui et al. (2011). For a recent review of all these developments we refer to Walther (2009).

To fix ideas, let p be a probability density on \mathbb{R} . We call p log-concave if it may be written as $p(x) = \exp \varphi(x)$ for some concave function $\varphi : \mathbb{R} \rightarrow [-\infty, \infty)$ and $x \in \mathbb{R}$. Based on a sample of i.i.d. random variables $V_1, \dots, V_n \in \mathbb{R}$ from p we seek to estimate this density via maximizing the normalized log-likelihood function

$$\ell(\varphi) = n^{-1} \sum_{i=1}^n \log p(V_i) = n^{-1} \sum_{i=1}^n \varphi(V_i)$$

over all concave functions $\varphi : \mathbb{R} \rightarrow [-\infty, \infty)$ such that $\int \exp \varphi(x) dx = 1$.

The merits of imposing log-concavity as a shape constraint have been described in detail in Balabdaoui et al. (2009), Cule et al. (2009), Walther (2009), and Dümbgen and Rufibach (2011). The most relevant of those properties in the current context are that many parametric models have log-concave densities. Examples include the Normal, Exponential, Uniform, Gamma(r, λ) for $r \geq 1$, Beta(a, b) for $a, b \geq 1$, generalized Pareto if the tail index γ is in $[-1, 0]$, Gumbel, Fréchet, Logistic or Laplace, to mention only some of these models. In addition, log-concavity can be considered a straightforward generalization of normality that shares many properties of the Normal distribution (Schuhmacher et al., 2011), but is much more flexible. For these reasons, assuming log-concavity offers a flexible nonparametric alternative to purely parametric models that includes a wide range of asymmetric, unimodal densities.

The crucial feature of \hat{p}_n , the log-concave density estimator and maximizer of $\ell(\varphi)$, for our intended application is that the corresponding CDF estimator \hat{P}_n is, under some regularity conditions and if the true density is in fact log-concave, asymptotically equivalent to \mathbb{P}_n . As a consequence, \hat{P}_n can be regarded as a smoother of the empirical distribution function \mathbb{P}_n for finite sample sizes n and this will directly translate into a smooth estimator for the ROC curve. Apart from the asymptotic result in Theorem A.1 given in the appendix, Dümbgen and Rufibach (2009, Corollary 2.5) have also derived additional features of \hat{P}_n for finite n that further support the plug-in strategy indicated above. Namely, that $\hat{P}_n(V_1) = 0$, $\hat{P}_n(V_n) = 1$, and that $\hat{P}_n(x) \in [\mathbb{P}_n(x) - n^{-1}, \mathbb{P}_n(x)]$ whenever x is a knot point of the piecewise linear function $\hat{\varphi}_n$.

To conclude the discussion of log-concave density estimation we would like to mention the kernel smoothed version \hat{p}_n^* of \hat{p}_n introduced in Dümbgen and Rufibach (2009, Section 3) and generalized in Chen and Samworth (2011). Closed formulas can be derived to compute \hat{p}_n^* and the corresponding CDF \hat{P}_n^* , see Dümbgen and Rufibach (2011). Since log-concavity is preserved under convolution, the smoothed estimate \hat{p}_n^* remains log-concave if the applied kernel has this property. In Dümbgen and Rufibach (2009), the normal kernel was used with a bandwidth chosen such that the variance of \hat{p}_n^* equals that of the sample V_1, \dots, V_n . As for \hat{p}_n , this makes \hat{p}_n^* a fully automatic estimator.

Note that \hat{p}_n^* is rather close to \hat{p}_n , see Chen and Samworth (2011, Proposition 3) for a corresponding quantitative result.

Testing for log-concavity. Typically, shape constraints are motivated through substantive considerations as above. However, researchers may want to formally assess the hypothesis of log-concavity. Hazelton (2011) adapts Silverman’s bandwidth test (Silverman, 1982) to test log-concavity of densities in any dimension and shows that it has healthy power in several scenarios. The test works by constructing a kernel estimate of the target density in which log-concavity is enforced by the use of a sufficiently large bandwidth. The test statistic is the minimal bandwidth for which log-concavity is achieved and the null distribution is generated via bootstrap.

In a more exploratory manner, log-concavity can be visually assessed by comparing such an estimate to, e.g., a kernel estimate of the same data.

Illustration of the log-concave density estimates. To illustrate the log-concave estimate, we show in Figure 1 density estimates applied to the log of the carbohydrate antigen 19-9 measurements from the pancreatic cancer serum biomarker study, for the pancreatitis patients (i.e. controls) only. This data was initially analyzed in Wieand et al. (1989) and re-assessed, among others, in Pepe (2003, Example 1.3.3), Cai and Moskowitz (2004, Section 6), Wan and Zhang (2007), or Zhou and Lin (2008) to name a few. For details on the dataset we refer to Section 5. In the left plot of Figure 1, the two log-concave density estimates as well as the standard R kernel estimate are displayed. It seems safe to assume a log-concave density for this data – the test by Hazelton (2011) yields a p -value of $p = 0.84$ (based on 9999 bootstrap samples). Note that the log-concave estimate is able to capture the data’s pronounced skewness. The fact that the log of the estimated density is piecewise linear, see Section A in the appendix, can be seen through the potential sharp kinks in the density, as it is the case in Figure 1. It is clear that such kinks are alleviated by the smoothed version \hat{p}_n^* . However, when looking at the CDFs on the right side of Figure 1, we see that the seemingly “nicer” fit of \hat{p}_n^* on the density level does not imply a significant difference for the estimated CDFs. The right picture further emphasizes that, as claimed above, the log-concave CDF estimator (whether kernel smoothed or not) can be considered to be a smoother version of the empirical CDF, a fact theoretically supported by Theorem A.1.

3 A new ROC curve estimator

To define our new ROC curve estimator we first compute log-concave distribution function estimates \hat{F}_m and \hat{G}_n based on the samples X_1, \dots, X_m and Y_1, \dots, Y_n . The estimates are then plugged into (1) to get

$$\hat{R}_{m,n}(t) := R(t; \hat{F}_m, \hat{G}_n) = 1 - \hat{G}_n(\hat{F}_m^{-1}(1 - t))$$

for $t \in (0, 1)$. The corresponding estimate based on the CDFs derived from the kernel smoothed log-concave estimates is denoted by $\hat{R}_{m,n}^*(t)$. Being a function of the integral of a piecewise exponential function, both estimators are smooth, i.e. they are at least

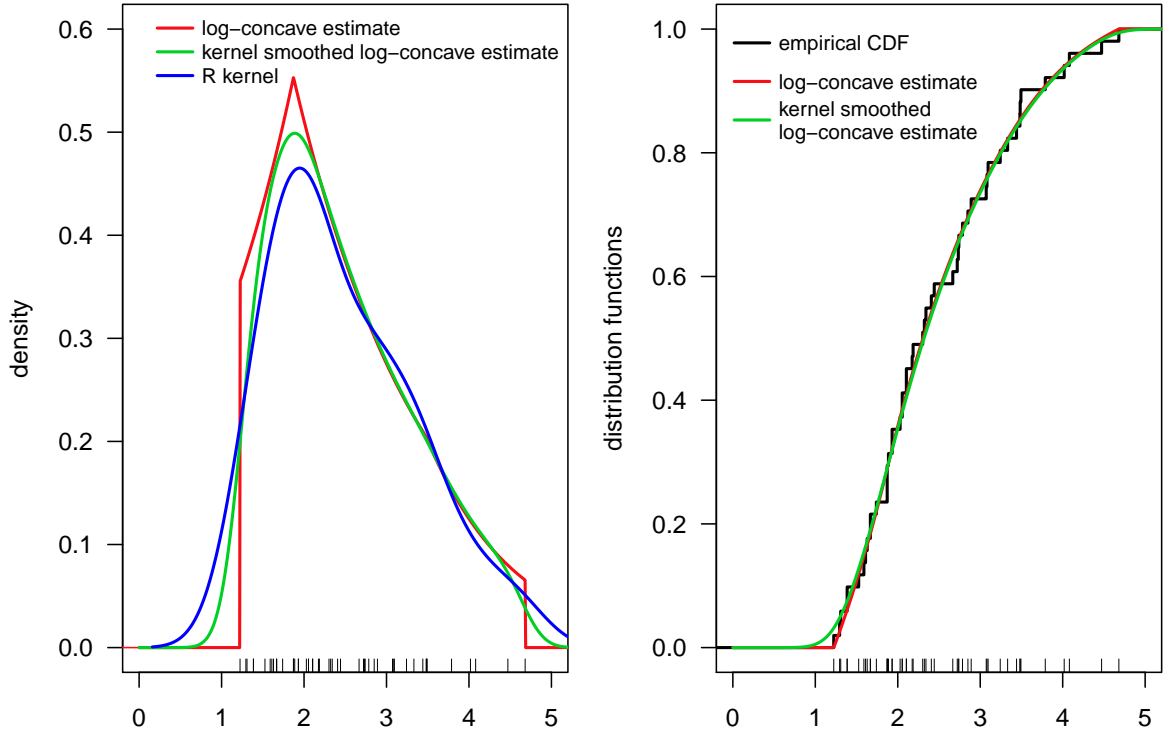


Figure 1: Pancreas dataset: Log-concave density estimation for log of CA 19.9 measurements, controls only.

differentiable on the entire domain $(0, 1)$ and infinitely often differentiable except on the joint set of knot points of \hat{f}_m and \hat{g}_n .

In Section A of the appendix we recapitulate asymptotic results for \hat{p}_n and \hat{P}_n . These results can be used to show that $\hat{R}_{m,n}$, as a process in t , is asymptotically equivalent to $\tilde{R}_{m,n}$ if the log-concavity assumption holds, see Section 4. This implies that the limiting behavior derived for $\tilde{R}_{m,n}$ in Hsieh and Turnbull (1996) in some sense applies to $\hat{R}_{m,n}$ or that asymptotic confidence intervals derived for $\tilde{R}_{m,n}(t)$ are also valid for $\hat{R}_{m,n}(t)$. In addition, $\hat{R}_{m,n}$ is smooth and, as we show in Section 6, for finite m and n typically more efficient than $\tilde{R}_{m,n}$. We therefore advocate the use of our new estimator as a surrogate for $\tilde{R}_{m,n}$ and an alternative to the kernel estimate $\hat{R}_{m,n}$ when it is safe to assume that the constituent densities are log-concave.

It was shown in Lloyd and Yong (1999) that the estimated ROC curve based on kernels outperforms the empirical, just as $\hat{R}_{m,n}$. One of the main advantages of shape constraint estimation in general, and in our current setting in particular, is that such estimates are fully automatic, i.e. they do not necessitate the choice of a kernel, bandwidth, or some other regularization parameter whose optimal value typically depends on the unknown function to be estimated.

Admittedly, and as it is typical for semiparametric ROC curve estimates, $\hat{R}_{m,n}$ is generally not invariant with respect to monotone transformations of either X and/or Y . However, this is the case for virtually all parametric models. A ROC curve R (true or estimated) is biased if there exists $p \in (0, 1)$ such that $R(p) < p$. Log-concave ROC curve estimates can indeed be biased. However, the bias is in general constrained to

regions of $(0, 1)$ that are small, and typically smaller compared to the binormal model. See Pepe (2003, p. 83) for a detailed discussion of potential bias in binormal models. Compared to a parametric model with Normal distributions for F and G , the log-concave approach is certainly more flexible, at the cost of only a small reduction in efficiency in case F and G are in fact Gaussian, as we verify in our simulation study in Section 6.

4 Main result

In the sequel, a function g is said to belong to the Hölder class $\mathcal{H}^{\beta,L}(I)$ of functions with exponent $\beta \in (1, 2]$ and constant $L > 0$ if for all $x, y \in I$ we have $|g'(x) - g'(y)| \leq L|x - y|^{\beta-1}$. The claim of asymptotic equivalence of $\hat{R}_{m,n}$ and $\tilde{R}_{m,n}$ under log-concavity is based on Theorem 4.1.

Theorem 4.1. *Assume that*

- C1 F is supported on $I_F = [0, 1]$ with positive density $F' = f$, and G has (potentially infinite) support $I_G \subseteq \mathbb{R}$.*
- C2 F and G both have log-concave densities $f = \exp \varphi, g = \exp \gamma$ where φ, γ are both Hölder-continuous with exponent $\beta \in (1, 2]$ and constant $L > 0$ on I_F and $[A, B] \subset I_G$, respectively.*
- C3 F and G have log-densities such that $\varphi'(x) - \varphi'(y) \geq C_F(x - y)$ for $0 \leq x < y \leq 1$ and $\gamma'(x) - \gamma'(y) \geq C_G(x - y)$ for $A \leq x < y \leq B$ and two constants $C_F, C_G > 0$.*
- C4 The sample size m is a function of n such that $n/m \rightarrow \lambda > 0$ as $n \rightarrow \infty$.*

Then, as $n \rightarrow \infty$,

$$\sqrt{n} \sup_{t \in J} \left(\hat{R}_{m,n}(t) - \tilde{R}_{m,n}(t) \right) \rightarrow_p 0$$

for $J = [F(A + \delta), F(B - \delta)]$ for every $\delta > 0$.

A strong approximation result for the empirical ROC curve process was provided in Hsieh and Turnbull (1996, Theorem 2.2), see also Pepe (2003, Result 5.2) or Horváth et al. (2008, Section 2), and implies the following corollary:

Corollary 4.2. *Let $\mathbb{B}_1(t)$ and $\mathbb{B}_2(t)$ be two independent Brownian Bridges. Then, as $n \rightarrow \infty$,*

$$\sqrt{n} \left(\hat{R}_{m,n}(t) - R(t) \right) \rightarrow_d \mathbb{B}_1 \left(1 - R(t) \right) + \sqrt{\lambda} \frac{g(F^{-1}(1 - t))}{f(F^{-1}(1 - t))} \mathbb{B}_2(1 - t)$$

uniformly on J .

Note that the original result in Hsieh and Turnbull (1996, Theorem 2.2) relating the empirical and the true ROC holds a.s. where Brownian Bridges depending on n are involved. This is the reason why we only get convergence in distribution in Corollary 4.2. The proof of and some comments on the assumptions of Theorem 4.1 can be found in Section D in the appendix.

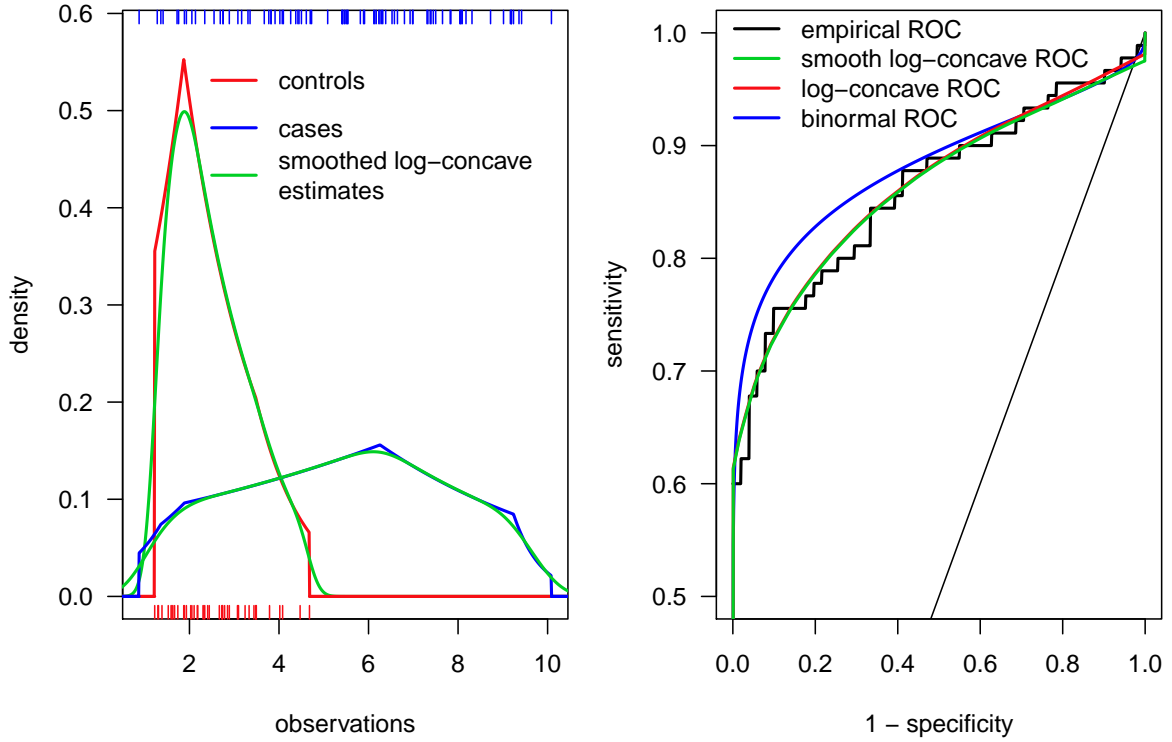


Figure 2: Pancreas dataset: Log-concave density estimates for log of CA 19.9 measurements and induced ROC curves.

5 Illustration of the new ROC curve estimate

We illustrate our new estimator and compare it to the empirical $\tilde{R}_{m,n}$ and binormal estimate $\bar{R}_{m,n}$ on the pancreatic cancer serum biomarker dataset. The study examined two biomarkers, an antigenic determinant denoted CA125 and a carbohydrate antigen denoted as CA19.9. The data consists of 51 controls that did not have cancer but suffered from pancreatitis and 90 pancreatic cancer patients. Here, we use the fully parametric binormal model, i.e. we estimate a and b in (2) directly from the mean μ and variance σ^2 of the underlying distributions. The results are displayed in Figure 2 and we observe that (1) Given that the empirical estimate follows the true curve closely, the binormal model does not seem to be appropriate here, (2) Both log-concave estimates are very close to the empirical estimate, acting as smoothers of the empirical estimator as discussed above, (3) The difference between the log-concave and the smoothed log-concave estimate in estimation of the ROC curve is small and mainly concentrated at the boundaries of the domain of the ROC curve. This is not surprising as the corresponding CDF estimates exhibit the same behavior as shown in Figure 1. This can be explained by the fact that the discontinuities of \hat{p}_n at the smallest and largest order statistic of the sample under consideration are smoothed out by \hat{p}_n^* . Finally, the test by Hazelton (2011) yields a p -value to assess the null hypothesis of log-concavity of $p = 0.84$ for the controls and $p = 0.83$ for the cases.

6 Simulations

Having shown asymptotic equivalence to the empirical estimate, it remains to empirically verify that, apart from its inherent properties such as smoothness and no need to choose any regularization parameters such as kernel and bandwidth, we additionally enjoy a gain in efficiency by using $\hat{R}_{m,n}$ instead of $\tilde{R}_{m,n}$. To this end, we have performed a simulation study for the scenarios provided in Table 1. Scenarios 1–3 serve as a benchmark for comparing our estimators to the binormal model. In addition, Scenario 2 has been analyzed in Lloyd and Yong (1999) and Scenario 3 in Zhou and Lin (2008), what enables a direct comparison to these estimators. Scenario 4 is intended to evaluate the effect of only one distribution being non-normal but skewed. This seems a plausible setup in ROC applications where values for the controls may be normal but those of cases are right-skewed. Scenario 5 assesses the performance in case both distributions are right-skewed. Finally, Scenario 6 evaluates the methods for symmetric but non-normal distributions.

Table 1: Scenarios we simulated from. $\text{Ga}(\alpha, \beta)$: Gamma distribution with shape parameter α and rate parameter β . $\text{Log}(\mu, s)$: Logistic distribution with location parameter μ and shape parameter s .

Scenario	F	m	G	n	has been used in
1	$N(0, 1)$	20	$N(1, 1)$	20	
2	$N(0, 1)$	100	$N(1, 1)$	100	Lloyd and Yong (1999, Figure 1)
3	$N(0, 1)$	100	$N(2, 1.2)$	100	Zhou and Lin (2008, Figure 1, left plot)
4	$N(2, 1)$	100	$\text{Ga}(2, 1)$	100	
5	$\text{Ga}(2, 1)$	100	$\text{Ga}(4, 1.5)$	100	
6	$\text{Log}(0, 1)$	100	$\text{Log}(2, 1)$	100	

In our simulations, we compare the empirical ROC estimate $\tilde{R}_{m,n}$, the ROC estimates based on the two variants of log-concave density estimate (maximum likelihood and its smoothed version), the kernel estimate $\hat{R}_{m,n}$ with optimal bandwidths to estimate the ROC curve introduced in Hall and Hyndman (2003), and the fully parametric binormal model. The latter model is well-specified only if both F and G have a normal density. We would like to emphasize that the binormal model was chosen as a competitor because it serves as a parametric benchmark if in fact F and G indeed have a normal density and our simulations show, that $\hat{R}_{m,n}$ and $\hat{R}_{m,n}^*$ are only slightly worse than this benchmark. However, we generally discourage the use of binormal models in practice.

To evaluate the performance of our competitors we use the average square error (ASE), defined for a ROC curve estimate \hat{R} of a true ROC curve R as

$$ASE(\hat{R}) = n_{\text{grid}}^{-1} \sum_{k=1}^{n_{\text{grid}}} \left(\hat{R}(u_k) - R(u_k) \right)^2$$

for grid points $u_i, i = 1, \dots, n_{\text{grid}}$. This criterion has been used in Zhou and Lin (2008) to evaluate different variants of binormal ROC curve estimates. Following the same approach as in the latter paper we choose the u_i 's to be equidistant on $[0, 1]$, $n_{\text{grid}} = 100$ and we generated $M = 500$ samples for each scenario. Using the same criterion and setting

all the simulation parameters to the values already used in (Zhou and Lin, 2008, left plot of Figure 1) enables direct comparison of our Scenario 3 to their results and we provide the analogous plot in Figure 4. In Figure 3 we display the results for all the six scenarios from Table 1, reported as follows: For each simulation run, the ASE is computed for the four estimators and $(ASE(\hat{R})_j/ASE(\tilde{R}_{m,n})_j)^{1/2}$ for $\hat{R} \in \{\hat{R}_{m,n}, \hat{R}_{m,n}^*, \underline{\hat{R}}_{m,n}, \bar{R}_{m,n}\}$ and $j = 1, \dots, M$ are reported. Thus, each of the four estimates is benchmarked with respect to the empirical ROC curve $\tilde{R}_{m,n}$. Figure 3 allows for the following observations: In the purely normal setting (Scenarios 1–3), the log-concave estimators are generally more efficient than the empirical, to an even remarkable extent for Scenario 3. We attribute this to the fact that here, the ROC curve is rather steep for small values on the abscissa, a shape that $\hat{R}_{m,n}$ and $\hat{R}_{m,n}^*$ are able to more efficiently capture than the empirical. The kernel estimator $\underline{\hat{R}}_{m,n}$ outperforms the empirical to an extent comparable to the log-concave estimates. Finally, compared to the fully parametric binormal model the loss of efficiency for the log-concave versions and the kernel estimate is notably small. When “visually averaging” the pointwise root mean square reductions over $t \in (0, 1)$ in Lloyd and Yong (1999, Figure 1, left plot) we see that their kernel estimate does approximately 15% better than the empirical estimate, for data simulated from our Scenario 2. This roughly corresponds to our median reduction in that scenario. The gain of the log-concave ROC estimates over the empirical estimate is maintained for non-normal though still log-concave setups (Scenarios 4–6), whereas in these cases the binormal model is certainly misspecified and therefore underperforming. In Scenarios 5 and 6, $\hat{R}_{m,n}$ and $\hat{R}_{m,n}^*$ provide values of $(ASE)^{1/2}$ with less variability than $\underline{\hat{R}}_{m,n}$. The effect of smoothing the maximum likelihood density estimate in ROC curve estimation is for all the analyzed scenarios rather small. This is to be expected given that these two estimates are close (Chen and Samworth, 2011, Proposition 3) and the fact that they are integrated when computing the ROC curve. In Zhou and Lin (2008) a new estimate for a and b in the semiparametric normal model is introduced and compared to a few competitors, e.g. the one from Cai and Moskowitz (2004). Note that all these estimators are rather complicated to compute and, to the best of our knowledge, no easy accessible implementation for these estimators is available. The left plot in Figure 1 of Zhou and Lin (2008) reports results for our Scenario 3 and the analogous result using the estimators from our simulation study is displayed in Figure 4. Although the estimates in Zhou and Lin (2008, Figure 1, left plot) were explicitly tailored for the binormal model, our log-concave estimates are only slightly, if at all, less efficient, which is consistent with the results shown in Figure 3. The same applies to $\underline{\hat{R}}_{m,n}$.

7 Bootstrap confidence intervals

To construct a confidence interval around $\hat{R}_{m,n}$ at a fixed point $t \in [0, 1]$, one can, by virtue of Theorem 4.1, use standard errors based on asymptotic theory for $\hat{R}_{m,n}$. Asymptotically, such an interval maintains the pre-specified coverage probability, provided that the log-concavity assumption holds. However, to exploit the gain in efficiency for finite samples in computation of confidence intervals for values of the true ROC curve, we suggest to proceed as sketched in Du and Tang (2009) for a similar spline estimator.

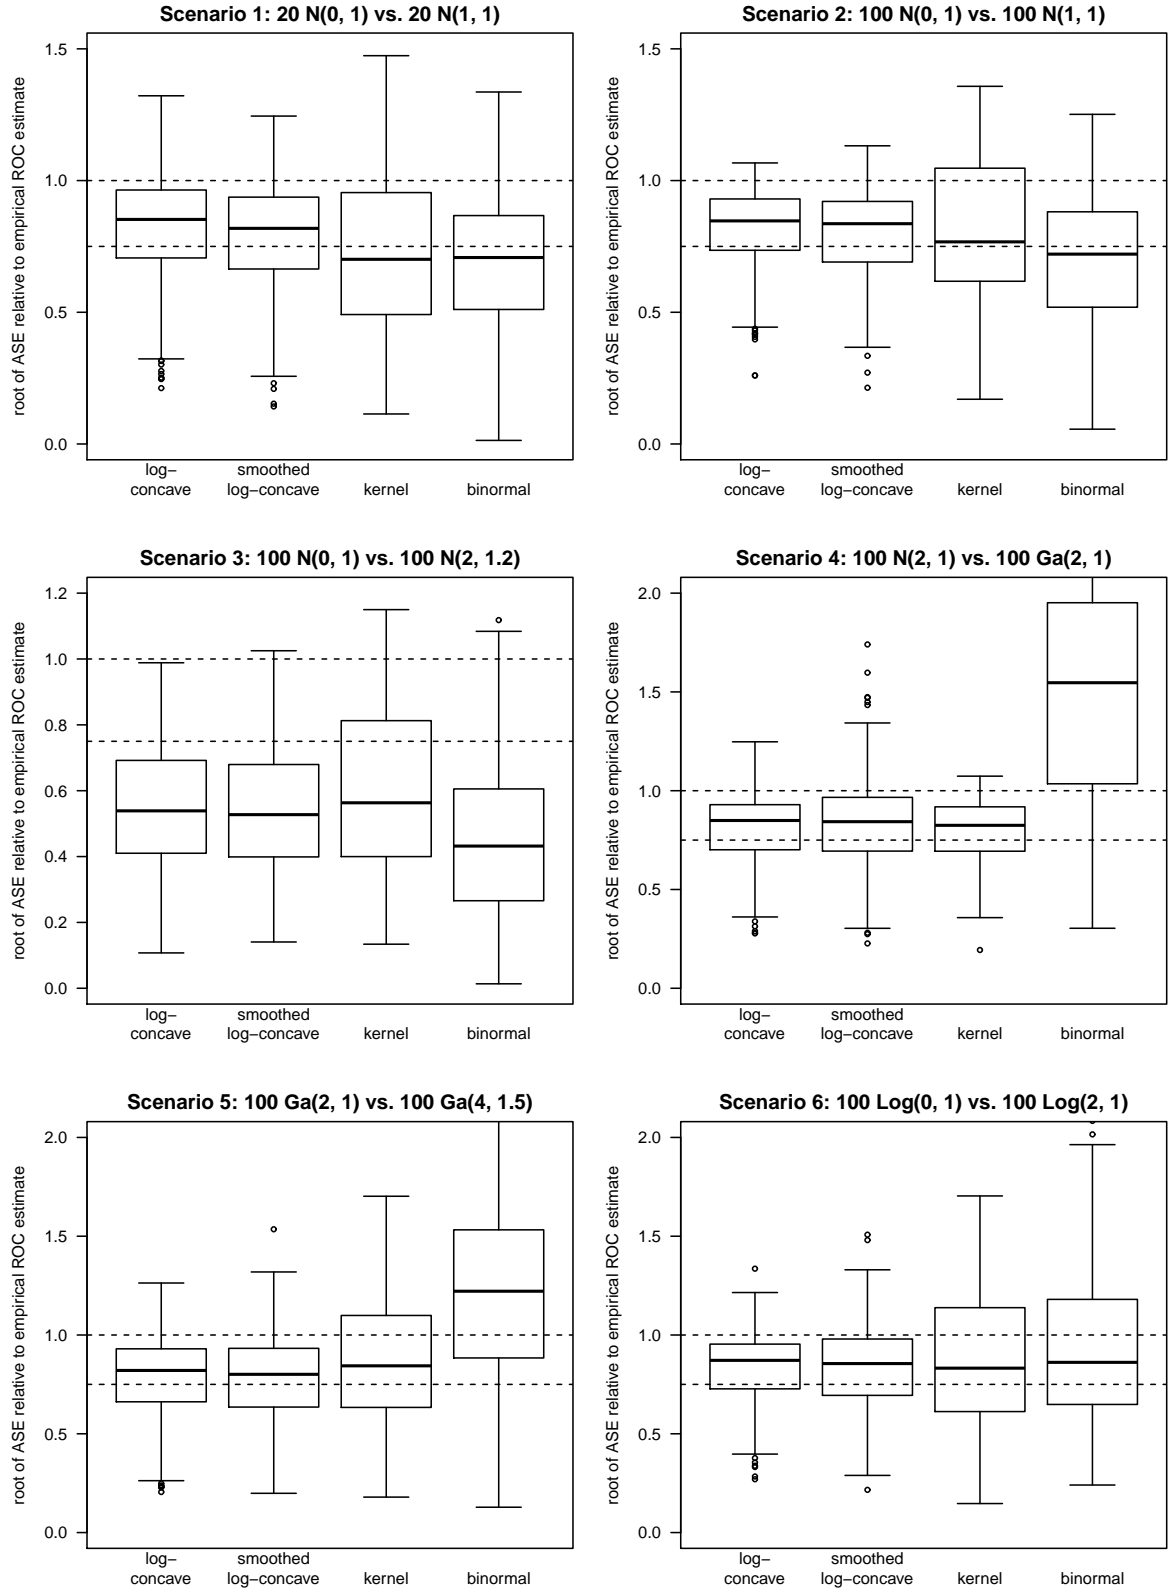


Figure 3: Results of simulation study. Horizontal dashed lines at 0.75, 1. Note the different scalings of the y -axis.

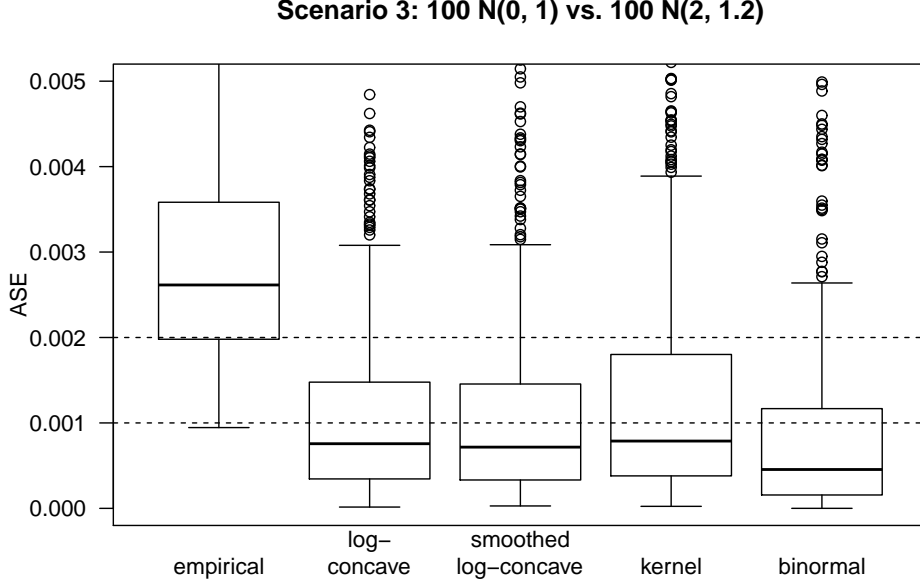


Figure 4: ASE for Scenario 3, to be compared with Zhou and Lin (2008, Figure 1, left plot). Horizontal dashed lines at 0.001, 0.002.

Namely, draw B bootstrap resamples $\{X_i^\#\}_{i=1}^m$ and $\{Y_i^\#\}_{i=1}^n$ from the original data and compute for each of these resamples the estimator $\hat{R}_{m,n}$. This yields bootstrapped versions $\hat{R}_{m,n}^\#$ of the estimated ROC curve. The $(1 - \alpha)$ confidence interval at a point $t \in [0, 1]$ can then be based on the quantiles of this bootstrap distribution, i.e. we compute $[(\hat{R}_{m,n}^\#(t))_{\alpha/2}, (\hat{R}_{m,n}^\#(t))_{1-\alpha/2}]$ where $(\hat{R}_{m,n}^\#(t))_{\alpha/2}$ and $(\hat{R}_{m,n}^\#(t))_{1-\alpha/2}$ are the $(\alpha/2)$ - and $(1 - \alpha/2)$ -quantile of the $\hat{R}_{m,n}^\#(t)$'s from the bootstrap samples.

To assess the performance of these intervals we proceeded as follows: For the scenarios in Table 1 we computed the above intervals at $t \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ for $B = 500$ and $1 - \alpha = 0.95$. As a benchmark we compare the proportion of these intervals out of $M = 500$ simulation runs that cover the true value $R(t)$ to the same proportion computed using the BTII interval proposed in Zhou and Qin (2005). BTII is a bootstrap confidence interval that does not only estimate the true positive fraction at a fixed t , but also accounts for the uncertainty due to the fact that we estimate the quantile corresponding to t based on the controls. Zhou and Qin (2005) show via simulations that BTII has superior coverage accuracy and shorter interval length compared to other approaches over a wide range of scenarios and for sample sizes comparable to those in our simulation setup.

In Figure 5 we present the results. We find that coverage proportions are basically identical for BTII and the interval based on $\hat{R}_{m,n}$, for $t \leq 0.7$. Zhou and Qin (2005) noted that the BTII interval performs poorly for values of t that correspond to high (≥ 0.95) sensitivities and small sample sizes. Looking at the smallest values of t where the true $R(t) \geq 0.95$ in our six scenarios from Table 1, we get values of 0.74/0.74/0.49/0.95/0.77/0.72, i.e. our results are in line with that rule-of-thumb: for Scenarios 1 and 3, BTII is

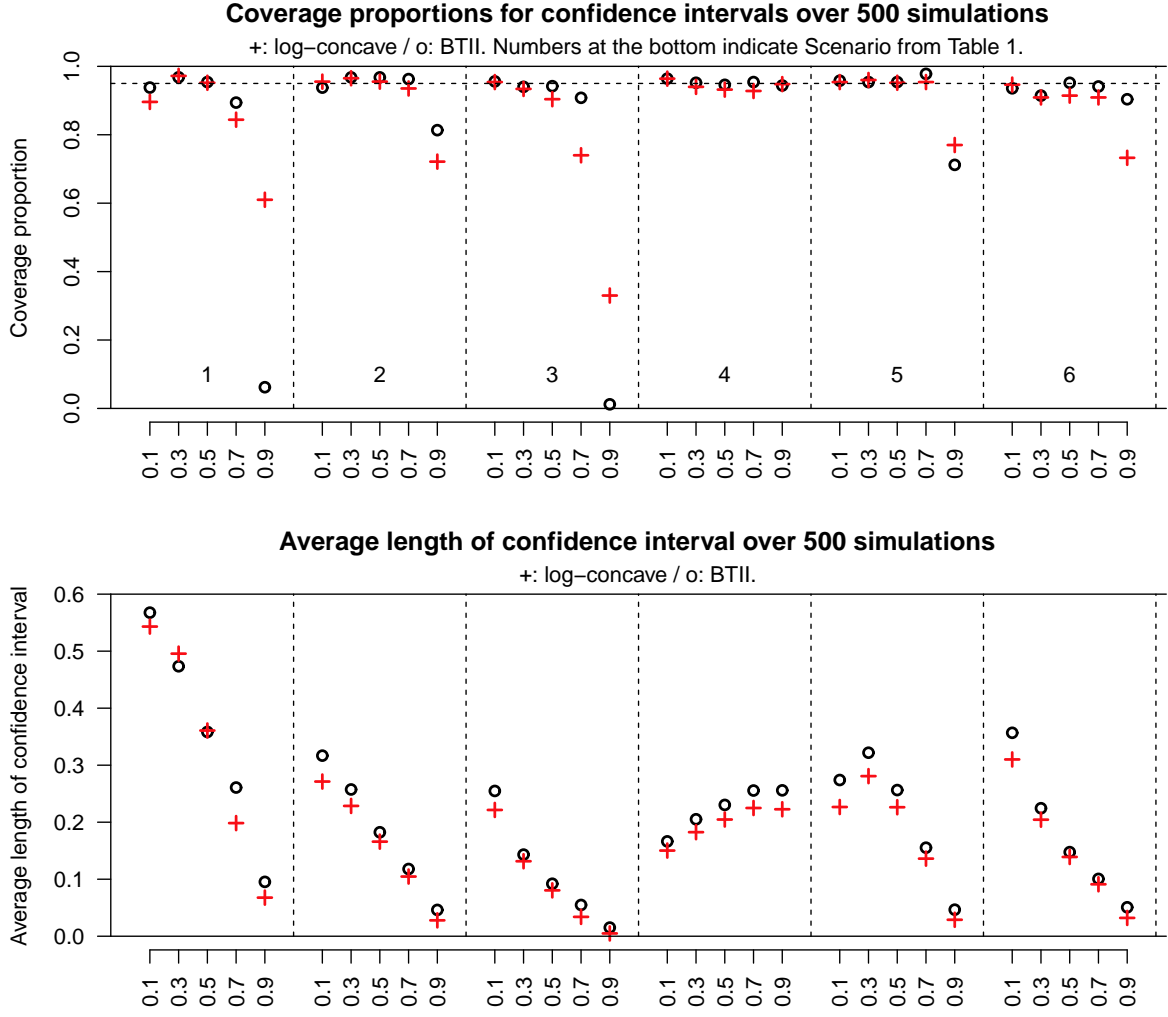


Figure 5: Performance of confidence intervals for R at a fixed t . Comparison of bootstrap intervals based on $\hat{R}_{m,n}$ and the BTII intervals discussed in Zhou and Qin (2005).

performing very poorly for $t = 0.9$, whereas the interval based on $\hat{R}_{m,n}$ has a better performance. Both methods approximately have the same difficulties in reaching the prescribed level for $t = 0.9$ in scenarios 2, 5, 6. However, as can be inferred from the lower plot in Figure 5, the intervals based on $\hat{R}_{m,n}$ generally yield confidence intervals with a shorter average length over the $M = 500$ simulations, a gain that is in line with the simulation results in Section 6.

8 Performance if f and/or g is not log-concave

First of all, we would like to point to recent work on estimation of a log-concave density under misspecification. Cule and Samworth (2010) show that the log-concave density estimator \hat{f}_n converges to the density g that minimizes the Kullback-Leibler distance to the true density f , in some exponentially weighted total variation norm. If f is

indeed log-concave, this implies a strong type of consistency for the log-concave density estimator.

Now, it is by no means immediate how this generalizes to our ROC curve estimate $\hat{R}_{m,n}$. To assess the robustness of $\hat{R}_{m,n}$, i.e. scenarios where at least one of the constituent distributions does not have a log-concave density, we extended the simulations reported in Section 6 to incorporate non-log-concave densities, see Section C in the appendix. Additional scenarios are given in Table 2. The conclusions from these simulations are that (1) $\hat{R}_{m,n}$ remains competitive even for moderate deviations from log-concavity for either f and/or g , and that (2) for small samples it even still outperforms the kernel-based estimator by Hall and Hyndman (2003).

9 Conclusions

We have presented a new approach to estimate ROC curves, which is based on plugging in distribution function estimates received from log-concave density estimates of the cases and controls data instead of their empirical counterparts. We propose bootstrap confidence intervals at a fixed t and show asymptotic equivalence between the empirical and the new ROC curve estimate if the log-concavity assumption holds. The performance of the new method is assessed via simulations and illustrated on a medical dataset. In what follows, we recapitulate the reasons why we consider our new approach useful.

In applications, it seems sensible to assume that the true underlying ROC curve R is in fact smooth, thus we expect a better performance of an estimate by imposing some degree of smoothness. In addition, it is visually appealing to have a smooth estimate rather than a jagged function like $\tilde{R}_{m,n}$.

Many parametric densities are in fact log-concave. As illustrated for the normal distribution in Scenarios 1 to 3 of our simulation study, not much efficiency is lost when considering log-concavity instead of the actual parametric model if the underlying distributions are indeed Normal. Besides, assuming log-concavity is much more flexible and robust to misspecification than imposing some parametric model, especially the widely used binormal model.

In our simulation study we have illustrated that over a wide spectrum of distributions, the estimators $\hat{R}_{m,n}$ and $\hat{R}_{m,n}^*$ perform remarkably better than the empirical ROC curve for finite sample sizes and have an efficiency that is comparable, and in some scenarios slightly better, than the kernel estimator $\hat{R}_{m,n}$. Confidence intervals based on $\hat{R}_{m,n}$ have comparable coverage probabilities but slightly shorter average interval lengths than the BTII interval from Zhou and Qin (2005), if the log-concavity assumption indeed holds. Taking into account these points, we advocate the use of our estimate since (1) we do not lose much, if anything at all, with respect to $\hat{R}_{m,n}$ or the binormal model and gain substantially with respect to the empirical ROC curve estimate but (2) are much more robust to misspecification than the binormal model. As opposed to kernel estimates, both our estimators $\hat{R}_{m,n}$ and $\hat{R}_{m,n}^*$ are fully automatic.

It was shown that, if the underlying densities are log-concave, $\hat{R}_{m,n}$ is asymptotically equivalent to $\tilde{R}_{m,n}$, so our new estimator can be considered a “smoother” of the empirical ROC curve. This “smoothing property” of \hat{P}_n has given rise to at least two

applications where empirical performance of a method could be substantially improved by using \hat{P}_n in place of \mathbb{P}_n as an estimator of the CDF. In Müller and Rufibach (2009) it was demonstrated that using quantiles based on \hat{P}_n^* instead of order statistics reduces estimation variability to a much greater extent in comparison to the bias introduced in tail index estimation in extreme value theory. This leads to a substantial reduction in mean squared error. Similarly, in Dümbgen and Rufibach (2011) a simulation study demonstrates increased power in the comparison of two distribution functions when using the Kolmogorov-Smirnov statistics based on the log-concave instead of the empirical distribution function estimate.

10 Software and reproducibility

Estimation of a univariate log-concave density and many additional related quantities is implemented in the R package **logcondens** (Dümbgen and Rufibach, 2011) available from CRAN. We have added the function `logConROC` to this package that gives the ROC curve estimates $\hat{R}_{m,n}$ and $\hat{R}_{m,n}^*$ as well as the corresponding area under the curve (AUC) in just one line of code. The bootstrap confidence intervals from Section 7 are implemented as the function `confIntBootLogConROC_t0` in **logcondens**. Note that the active-set algorithm used to maximize the log-likelihood function for the log-concave density estimate is remarkably efficient so that it only takes seconds to compute $\hat{R}_{m,n}$ for rather large sample sizes m and n , e.g. of order $m = n = 10^4$. For small to moderate sample sizes computation is immediate. The **pancreas** dataset used to illustrate our estimate is also included in **logcondens** and thus readily accessible. The code to compute the kernel estimator by Hall and Hyndman (2003) is available from the author’s webpage. All this enables straightforward reproducibility of the plots and simulations. In addition, the function `smooth.roc` in the R package **pROC** (Robin et al., 2011) offers the option to fit a ROC curve based on log-concave density estimates.

This document was created using **Sweave** (Leisch, 2002), \LaTeX (Knuth, 1984; Lamport, 1994), R 2.14.0 (R Development Core Team, 2010) with the R packages **logcondens** (Dümbgen and Rufibach, 2011, Rufibach and Dümbgen, 2011, Version 2.0.6), **xtable** (Dahl, 2009, Version 1.6-0), and **cacheSweave** (Peng, 2008, Version 0.6).

Appendix

A Additional details on log-concave density estimation

Here, we provide additional details and references on log-concave density estimation. Dümbgen and Rufibach (2009) show that the maximizer $\hat{\varphi}_n$ of ℓ is unique, piecewise linear on the interval $[V_{(1)}, V_{(n)}]$ with knots only at (some of the) observations $V_{(i)}$, and $\hat{\varphi}_n = -\infty$ elsewhere. Here $V_{(1)} \leq V_{(2)} \leq \dots \leq V_{(n)}$ are the ordered observations, and a “knot” of $\hat{\varphi}_n$ is a location where this function changes slope. The MLEs $\hat{\varphi}_n$, $\hat{p}_n = \exp \hat{\varphi}_n$ and $\hat{P}_n(t) = \int_{-\infty}^t \hat{p}_n(x) dx$ are consistent with certain rates of convergence,

see Dümbgen and Rufibach (2009) and Balabdaoui et al. (2009) as well as Theorem A.1 below.

The crucial feature of the estimate in our context is summarized in Theorem A.1.

Theorem A.1. *[Dümbgen and Rufibach, 2009, Theorems 4.1 and 4.4] Assume that the log-density $\varphi = \log p \in \mathcal{H}^{\beta,L}(T)$ for some exponent $\beta \in [1, 2]$, some constant $L > 0$ and a subinterval $T = [A, B]$ of the interior of the support $\{p > 0\}$. Then,*

$$\max_{t \in T(n,\beta)} |\hat{\varphi}_n - \varphi|(t) = O_p(\rho_n^{\beta/(2\beta+1)}),$$

where $T(n, \beta) := [A + \rho_n^{1/(2\beta+1)}, B - \rho_n^{1/(2\beta+1)}]$ and $\rho_n = \log(n)/n$. In particular, this entails that for $\beta \in (1, 2]$

$$\max_{t \in T(n,\beta)} |\hat{P}_n(t) - \mathbb{P}_n(t)| = o_p(n^{-1/2})$$

where \mathbb{P}_n is the empirical CDF of V_1, \dots, V_n .

Note that the result for $\hat{\varphi}_n$ remains true when we replace $\hat{\varphi}_n - \varphi$ by $\hat{p}_n - p$. Furthermore, it is well-known that the rates of convergence in Theorem A.1 are optimal, even if β was known (Has'minskiĭ, 1978). Thus the log-concave (log-)density estimator adapts to the unknown smoothness of p in the range $\beta \in (1, 2]$.

B Estimation of the AUC

In Section 6 we have shown that $\hat{R}_{m,n}$ and $\hat{R}_{m,n}^*$ (and $\hat{R}_{m,n}$) are valuable approaches to estimate the ROC curve. In this section, we will show that this holds true also in estimation of the AUC. AUC is the most widely used (Pepe, 2003, Section 4.3.1) summary measure for ROC curves and typically reported in the analysis of a diagnostic test. First, note that Theorem 4.1 together with Hsieh and Turnbull (1996, Theorem 2.3) implies that

$$\begin{aligned} \sqrt{n}(\widehat{AUC} - AUC) &:= \sqrt{n} \int_J (\hat{R}_{m,n}(t) - R(t)) dt \\ &\rightarrow_d N(0, \sigma_{AUC}^2), \end{aligned}$$

so that the AUC based on $\hat{R}_{m,n}$ and $\tilde{R}_{m,n}$ share the same distributional limit (under the assumptions of Theorem 4.1) as well and asymptotic results for the empirical AUC are valid for the AUC based on the log-concave ROC. The expression for the asymptotic variance σ_{AUC}^2 of the empirical AUC is provided in Hsieh and Turnbull (1996, Theorem 2.3), see also the discussion in Pepe (2003, Section 5.2.5).

C Estimation of $\hat{R}_{m,n}$ and AUC under misspecification

Here, we report on additional simulations when either f and/or g are misspecified, i.e. not log-concave. We have chosen the Lomax densities as these are an often-used and

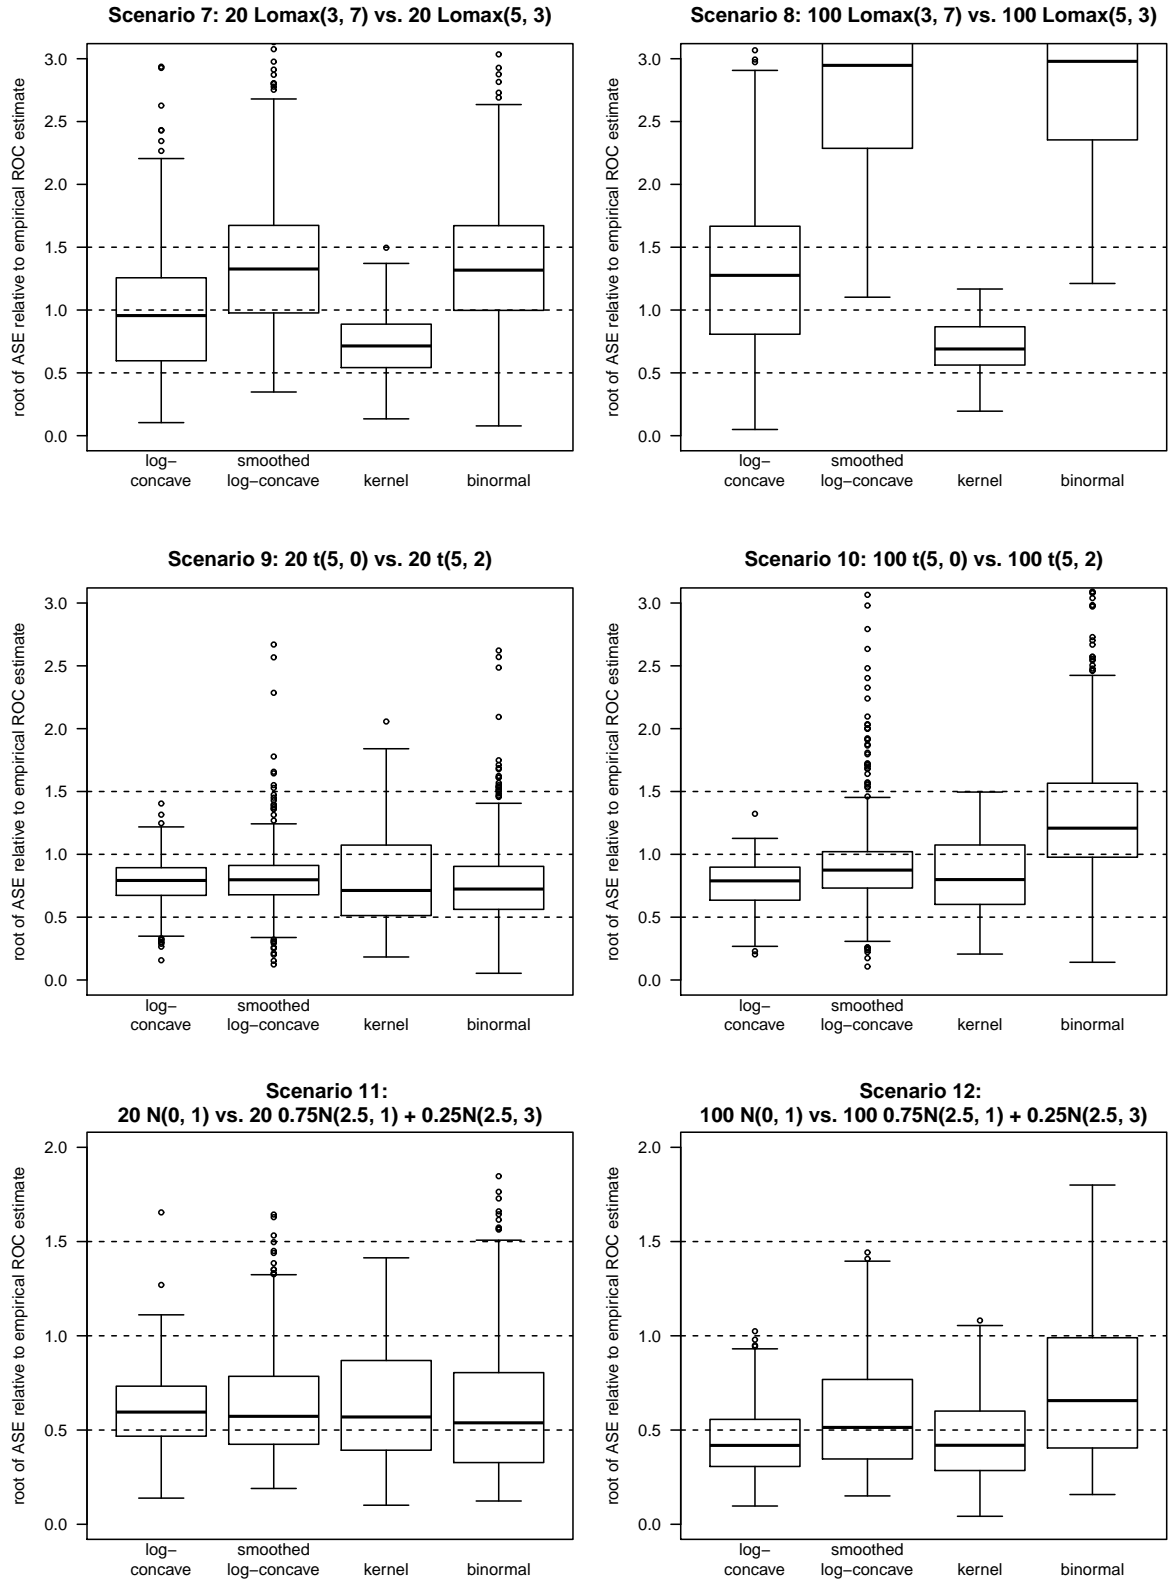


Figure 6: Results of simulation study for misspecified models. Horizontal dashed lines at 0.5, 1, 1.5. Note the different scalings of the y -axis.

Table 2: Scenarios we simulated from to assess performance of the new estimator in misspecified models.

Scenario	F	m	G	n
7	Lomax(3, 7)	20	Lomax(5, 3)	20
8	Lomax(3, 7)	100	Lomax(5, 3)	100
9	t(5, 0)	20	t(5, 2)	20
10	t(5, 0)	100	t(5, 2)	100
11	N(0, 1)	20	0.75·N(2.5, 1) + 0.25·N(2.5, 3)	20
12	N(0, 1)	100	0.75·N(2.5, 1) + 0.25·N(2.5, 3)	100

unbiased model. The t - and the normal mixture densities were selected as they are still unimodal but have slightly “too heavy” tails to still be log-concave.

The ASE relative to the empirical estimate for the simulations from Table 2 are displayed in Figure 6. As expected, the estimator by Hall and Hyndman (2003) now clearly outperforms $\hat{R}_{m,n}$. However, the advantage of the latter estimator is only substantial for the Lomax scenarios whereas for the t and normal mixture setups the two estimators $\hat{R}_{m,n}$ and $\hat{R}_{m,n}$ perform comparably, even with some advantage for $\hat{R}_{m,n}$. We thus conclude that $\hat{R}_{m,n}$ remains competitive even for moderate deviations from log-concavity in either F and/or G , and that for small samples it still outperforms the kernel-based estimator by Hall and Hyndman (2003).

As a matter of fact, if either F and/or G has a bimodal density it does not seem sensible to use $\hat{R}_{m,n}$ and we have thus omitted such scenarios from our misspecification simulation study. It seems safe to assume that in such a scenario, $\hat{R}_{m,n}$ will outperform $\hat{R}_{m,n}$ and that the latter should not be used based on substantive arguments.

As a final remark on misspecified models we would like to emphasize that one typically applies shape constraint estimates based on *substantive knowledge*, i.e. *knows* that the densities of F and G are log-concave, or one assesses this claim via formal testing, or via contrasting the log-concave to kernel density estimates. In addition, as it comes out of the simulations, $\hat{R}_{m,n}$ is somewhat robust against misspecification.

D Proof of Theorem 4.1

We write

$$\begin{aligned}
& \sqrt{n} \sup_{t \in J} \left| R(t; \mathbb{F}_m, \mathbb{G}_n) - R(t; \hat{F}_m, \hat{G}_n) \right| = \\
&= \sqrt{n} \sup_{t \in J} \left| (1 - \mathbb{G}_n(\mathbb{F}_m^{-1}(1-t))) - (1 - \hat{G}_n(\hat{F}_m^{-1}(1-t))) \right| \\
&= \sqrt{n} \sup_{t \in J} \left| \hat{G}_n(\hat{F}_m^{-1}(t)) - \mathbb{G}_n(\mathbb{F}_m^{-1}(t)) \right| \\
&= \sqrt{n} \sup_{t \in J} \left| \hat{G}_n(\hat{q}_{1,t}) - \mathbb{G}_n(\hat{q}_{1,t}) + \hat{g}_n\left(\hat{q}_{1,t} + \theta_m(\hat{q}_{2,t} - \hat{q}_{1,t})\right)(\hat{q}_{2,t} - \hat{q}_{1,t}) \right| \\
&\leq \sqrt{n} \sup_{t \in J} \left| \hat{G}_n(\hat{q}_{1,t}) - \mathbb{G}_n(\hat{q}_{1,t}) \right| + \sqrt{n} \sup_{t \in J} \left| \hat{g}_n\left(\hat{q}_{1,t} + \theta_m(\hat{q}_{2,t} - \hat{q}_{1,t})\right)(\hat{q}_{2,t} - \hat{q}_{1,t}) \right| \\
&=: T_1(n, m) + T_2(n, m)
\end{aligned}$$

where we defined $\hat{q}_{1,t} = \mathbb{F}_m^{-1}(t)$ and $\hat{q}_{2,t} = \hat{F}_m^{-1}(t)$ and θ_m is some value in $(0, 1)$ for all m . By the Bahadur-Kiefer theorem (see e.g. van der Vaart and Wellner, 1996, Example 3.9.24 or van der Vaart, 1998, Section 21.2) we can write for any $t \in J$ $\mathbb{F}_m^{-1}(F(t)) = t + (\mathbb{F}_m(t) - F(t)) + o_{\text{as}}(n^{-1/2})$, what together with strong consistency of \mathbb{F}_m (van der Vaart, 1998, p. 265) implies that for any $\varepsilon \in (0, \delta]$ we can find an n_0 large enough so that for all $m, n \geq n_0$ almost surely

$$\left| \mathbb{F}_m^{-1}(F(A + \delta)) - A - \delta - \rho_n \right| \leq \varepsilon$$

where $\rho_n = \log(n)/n$. But this implies that $\mathbb{F}_m^{-1}(F(A + \delta)) \geq A + \rho_n$ and similarly $\mathbb{F}_m^{-1}(F(B - \delta)) \leq B - \rho_n$ with probability one. Thus,

$$\begin{aligned} T_1(n, m) &= \sqrt{n} \sup_{t \in [F(A+\delta), F(B-\delta)]} \left| \hat{G}_n(\hat{q}_{1,t}) - \mathbb{G}_n(\hat{q}_{1,t}) \right| \\ &= \sqrt{n} \sup_{u \in [\mathbb{F}_m^{-1} \circ F(A+\delta), \mathbb{F}_m^{-1} \circ F(B-\delta)]} \left| \hat{G}_n(u) - \mathbb{G}_n(u) \right| \end{aligned}$$

and the latter expression is (almost surely) not smaller than

$$\sqrt{n} \sup_{u \in [A+\rho_n, B-\rho_n]} \left| \hat{G}_n(u) - \mathbb{G}_n(u) \right| = o_p(1)$$

by virtue of the second statement of Theorem 4.4 in Dümbgen and Rufibach (2009).

As for $T_2(m, n)$ we first note \hat{g}_n is bounded with probability one (Pal et al., 2007, Theorem 3.2 and Cule and Samworth, 2010, Lemma 3). Then, using again Dümbgen and Rufibach (2009, Theorem 4.4) one can show that under the assumptions C2 and C3

$$\sup_{t \in J} |\hat{F}_m(t) - F(t)| \leq \sup_{t \in J} |\mathbb{F}_m(t) - F(t)| + o_p(n^{-1/2}) \quad (3)$$

what implies that $\sqrt{n}(\hat{F}_m - F)$ converges weakly in $D[J]$ (the space of cadlag functions on J) to the same limit as $\sqrt{n}(\mathbb{F}_m - F)$, namely to $\mathbb{B} \circ F$ for a standard Brownian Bridge \mathbb{B} . But convergence of the estimated distribution function in $D[J]$ implies convergence of the corresponding quantile process in $D[J]$, by virtue of van der Vaart and Wellner (1996, Lemma 3.9.23). Thus, for some generic constant $C > 0$,

$$\begin{aligned} T_2(n, m) &= \sqrt{n} \sup_{t \in J} \left| \hat{g}_n \left(\hat{q}_{1,t} + \theta_m (\hat{q}_{2,t} - \hat{q}_{1,t}) \right) (\hat{q}_{2,t} - \hat{q}_{1,t}) \right| \\ &\leq C \sqrt{\lambda} (1 + o_p(1)) \sqrt{m} \sup_{t \in J} \left| \hat{F}_m^{-1}(t) - \mathbb{F}_m^{-1}(t) \right| \\ &\leq C \sqrt{\lambda} (1 + o_p(1)) \sqrt{m} \sup_{t \in J} \left\{ \left| \hat{F}_m^{-1}(t) - F^{-1}(t) + \frac{\mathbb{B} \circ F(F^{-1}(t))}{f(F^{-1}(t))} \right| + \right. \\ &\quad \left. \left| \mathbb{F}_m^{-1}(t) - F^{-1}(t) + \frac{\mathbb{B} \circ F(F^{-1}(t))}{f(F^{-1}(t))} \right| \right\} = o_p(1) \end{aligned}$$

via van der Vaart and Wellner (1996, Lemma 3.9.23). \square

For the empirical ROC curve this is carried out in Tang et al. (2008, Web supplement) and Tang and Zhou (2009, Web supplement).

Now, by Theorem 4.3 in Dümbgen and Rufibach (2009), again under the assumptions of Theorem 4.1, we have that $\sqrt{n}(\hat{F} - F)$ has the same limiting distribution as $\sqrt{n}(\mathbb{F} - F)$. This in turn implies that the above result carries over to $\hat{R}_{m,n}$.

Comments on the assumptions of Theorem 4.1. As it can be inferred from the proof, the crucial assumption is not log-concavity of the underlying densities, but rather the Dvoretzky-Kiefer-Wolfowitz inequality (3) which must hold for the CDF estimate. Indeed, log-concavity is only one possible assumption that entails this property. The density whose quantile function is involved in the ROC curve estimate needs to have finite support, since it must fulfill the assumptions for the Bahadur-Kiefer theorem (van der Vaart, 1998, Section 21.2). However, it is easy to see that any truncated log-concave density remains log-concave, so that this assumption does not seem to be too restrictive.

References

- BALABDAOUI, F., JANKOWSKI, H. and RUFIBACH, K. (2011). Maximum likelihood estimation and confidence bands for a discrete log-concave distribution. *ArXiv e-prints* .
URL <http://arxiv.org/abs/1107.3904>
- BALABDAOUI, F., RUFIBACH, K. and WELLNER, J. A. (2009). Limit distribution theory for maximum likelihood estimation of a log-concave density. *Ann. Statist.* **37** 1299–1331.
- CAI, T. and MOSKOWITZ, C. S. (2004). Semi-parametric estimation of the binormal roc curve for a continuous diagnostic test. *Biostatistics* **5** 573–586.
- CHEN, Y. and SAMWORTH, R. J. (2011). Smoothed log-concave maximum likelihood estimation with applications. *ArXiv e-prints* .
- CULE, M., GRAMACY, R. and SAMWORTH, R. (2009). LogConcDEAD: An R package for maximum likelihood estimation of a multivariate log-concave density. *Journal of Statistical Software* **29**.
- CULE, M. and SAMWORTH, R. (2010). Theoretical properties of the log-concave maximum likelihood estimator of a multidimensional density. *Electronic J. Stat.* **4** 254–270.
- CULE, M., SAMWORTH, R. and STEWART, M. (2010). Maximum likelihood estimation of a multidimensional log-concave density. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **72** 545–607.
- DAHL, D. B. (2009). *xtable: Export tables to LaTeX or HTML*. R package version 1.5-6.
- DU, P. and TANG, L. (2009). Transformation-invariant and nonparametric monotone smooth estimation of ROC curves. *Stat. Med.* **28** 349–359.
- DÜMBGEN, L., HÜSLER, A. and RUFIBACH, K. (2010). Active set and EM algorithms for log-concave densities based on complete and censored data. Tech. rep., University of Bern. Available at [arXiv:0707.4643](https://arxiv.org/abs/0707.4643).

- DÜMBGEN, L. and RUFIBACH, K. (2009). Maximum likelihood estimation of a log-concave density and its distribution function. *Bernoulli* **15** 40–68.
- DÜMBGEN, L. and RUFIBACH, K. (2011). **logcondens**: Computations related to univariate log-concave density estimation. *Journal of Statistical Software* **39** 1–28.
- HALL, P. G. and HYNDMAN, R. J. (2003). Improved methods for bandwidth selection when estimating ROC curves. *Statist. Probab. Lett.* **64** 181–189.
- HARRELL, F. E. and DAVIS, C. E. (1982). A new distribution-free quantile estimator. *Biometrika* **69** 635–640.
- HAS'MINSKIĬ, R. Z. (1978). A lower bound for risks of nonparametric density estimates in the uniform metric. *Teor. Veroyatnost. i Primenen.* **23** 824–828.
- HAZELTON, M. L. (2011). Assessing log-concavity of multivariate densities. *Statist. Probab. Lett.* **81** 121–125.
- HORVÁTH, L., HORVÁTH, Z. and ZHOU, W. (2008). Confidence bands for ROC curves. *J. Statist. Plann. Inference* **138** 1894–1904.
- HSIEH, F. and TURNBULL, B. W. (1996). Nonparametric and semiparametric estimation of the receiver operating characteristic curve. *Ann. Statist.* **24** 25–40.
- KNUTH, D. (1984). Literate programming. *Computer Journal* **27** 97–111.
- LAMPORT, L. (1994). *TEX: A Document Preparation System*. 2nd ed. Addison-Wesley, Reading, Massachusetts.
- LEISCH, F. (2002). Dynamic generation of statistical reports using literate data analysis. In *COMPSTAT 2002 – Proceedings in Computational Statistics* (W. Härdle and B. Rönz, eds.). Physica Verlag, Heidelberg.
- LLOYD, C. (1998). Using smoothed receiver operating characteristic curves to summarize and compare diagnostic systems. *J. Amer. Statist. Assoc.* **93** 1356–1364.
- LLOYD, C. J. (2002). Estimation of a convex ROC curve. *Statist. Probab. Lett.* **59** 99–111.
- LLOYD, C. J. and YONG, Z. (1999). Kernel estimators of the ROC curve are better than empirical. *Statist. Probab. Lett.* **44** 221–228.
- MÜLLER, S. and RUFIBACH, K. (2009). Smooth tail index estimation. *J. Statist. Comput. Simul.* **79** 1155–1167.
- PAL, J. K., WOODROOFE, M. B. and MEYER, M. C. (2007). Estimating a polya frequency function. In *Complex Datasets and Inverse Problems: Tomography, Networks, and Beyond*, vol. 54 of *IMS Lecture Notes-Monograph Series*. IMS, 239–249.
- PENG, L. and ZHOU, X.-H. (2004). Local linear smoothing of receiver operating characteristic (ROC) curves. *J. Statist. Plann. Inference* **118** 129–143.

- PENG, R. D. (2008). Caching and distributing statistical analyses in R. *Journal of Statistical Software* **26** 1–24.
- PEPE, M. S. (2003). *The statistical evaluation of medical tests for classification and prediction*, vol. 28 of *Oxford Statistical Science Series*. Oxford University Press, Oxford.
- QIU, P. and LE, C. (2001). ROC curve estimation based on local smoothing. *J. Statist. Comput. Simulation* **70** 55–69.
- R DEVELOPMENT CORE TEAM (2010). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
- ROBIN, X., TURCK, N., HAINARD, A., TIBERTI, N., LISACEK, F., SANCHEZ, J. C. and MULLER, M. (2011). pROC: an open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics* **12** 77.
- RUFIBACH, K. (2006). *Log-concave density estimation and bump hunting for I.I.D. observations*. Ph.D. thesis, Universities of Bern and Göttingen.
- RUFIBACH, K. (2007). Computing maximum likelihood estimators of a log-concave density function. *J. Statist. Comp. Sim.* **77** 561–574.
- RUFIBACH, K. and DÜMBGEN, L. (2011). *logcondens: Estimate a Log-Concave Probability Density from iid Observations*. R package version 2.0.5.
- SCHUHMACHER, D., HÜSLER, A. and DÜMBGEN, L. (2011). Multivariate log-concave distributions as a nearly parametric model. *Statistics & Decisions* **28** 1001–1017.
- SILVERMAN, B. W. (1982). On the estimation of a probability density function by the maximum penalized likelihood method. *Ann. Statist.* **10** 795–810.
- TANG, L., EMERSON, S. S. and ZHOU, X.-H. (2008). Nonparametric and semiparametric group sequential methods for comparing accuracy of diagnostic tests. *Biometrics* **64** 1137–1145.
- TANG, L. and ZHOU, X.-H. (2009). Semiparametric inferential procedures for comparing multivariate ROC curves with interaction terms. *Statist. Sinica* **19** 1203–1221.
- VAN DER VAART, A. W. (1998). *Asymptotic statistics*, vol. 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge.
- VAN DER VAART, A. W. and WELLNER, J. A. (1996). *Weak Convergence and Empirical Processes with Applications in Statistics*. Springer, New York.
- WALTHER, G. (2002). Detecting the presence of mixing with multiscale maximum likelihood. *J. Amer. Statist. Assoc.* **97** 508–513.
- WALTHER, G. (2009). Inference and modeling with log-concave distributions. *Statist. Sci.* **24** 319–327.

- WAN, S. and ZHANG, B. (2007). Smooth semiparametric receiver operating characteristic curves for continuous diagnostic tests. *Stat. Med.* **26** 2565–2586.
- WIEAND, S., GAIL, M. H., JAMES, B. R. and JAMES, K. L. (1989). A family of nonparametric statistics for comparing diagnostic markers with paired or unpaired data. *Biometrika* **76** 585–592.
- ZHOU, X. and HAREZLAK, J. (2002). Comparison of bandwidth selection methods for kernel smoothing of ROC curves. *Statist. Med.* **21** 2045–2055.
- ZHOU, X.-H. and LIN, H. (2008). Semi-parametric maximum likelihood estimates for ROC curves of continuous-scales tests. *Stat. Med.* **27** 5271–5290.
- ZHOU, X.-H. and QIN, G. (2005). A new confidence interval for the difference between two binomial proportions of paired data. *J. Statist. Plann. Inference* **128** 527–542.